

SNIC Emerging Technologies

Lustre v2+ facilities Changelogs and RobinHood

Author: Mathias Lindberg
Email: mathlin@chalmers.se
Date: 2016-06-30

Objective

Since Lustre in its current version is limited to reporting and enforcing user and group quotas only, additional components are needed for reporting project quotas.

Components needed

A Lustre version 2.1 change-log capable filesystem and RobinHood (<https://github.com/cea-hpc/robinhood>).

Implementation

Turn on Lustre-changelogs on your MDS. Register a reader on the MDS, and setup RobinHood to consume change-log entries with the reader id you just registered.

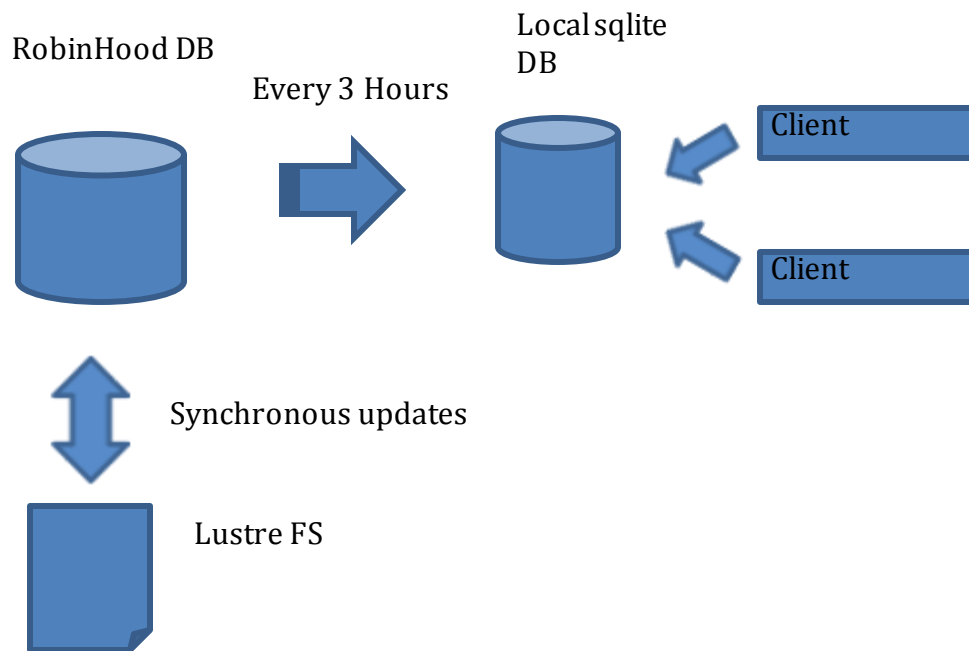
Note that once you enable change-logs on your MDS, RobinHood need to be running in order to consume and purge consumed records. Failure to do so can result in “unexpected behavior” on the MDS according to the documentation.

Start RobinHood with *scan* and *log reader* modules active. This will scan your entire filesystem and populate the database with information. While the log reader module will detect any immediate changes being made to the filesystem and stores that in the database.

C3SE additions

Since the database response times for a quota report and when a user interactively asked for quota information without special considerations on the i/o subsystem might take a while, we create a cached sqlite flat file database that contains user group and project quotas. This database is created by combining the information from RobinHood with information from the C3SE project and account database system Tengil that holds the user, group and project quota limits.

In order to ensure the continuous consumption and purging a change log records, a script that checks the index of the latest record and the latest consumed record. The script issues a warning if the difference exceeds configured limit.



Details

Pre-generated RPMs can be downloaded on Sourceforge, for the following configurations:

- x86_64 architecture, RedHat 5/6 Linux families
- MySQL database 5.x
- Posix filesystems, Lustre 1.8, 2.1, 2.2, 2.3, 2.4, 2.5, 2.6

Install the robinhood version that corresponds to the Lustre version you are running on you servers (or compile from source), the RPM *robinhood-tmpfs* contains the robinhood daemon, reporting tools and configuration templates. The RPM *robinhood-adm* contains tools to configure the DB host and Lustre MDS.

Use the *robinhood-adm* tools to create a database on your DB host and xxx to activate changelogs on your Lustre MDS.

Create a configuration file from the templates, make sure you enter the reader you registered on your MDS in the configuration file.

Start robinhood with the options *--scan --reader*, *--scan* is only needed the first time you run robinhood or if you have lost filesystem updates for any reason.

Default start options for robinhood is available in the */etc/sysconfig/robinhood* file.

Future additions

Current installation only warns the users, groups or project area responsible about the quota infraction. It would be preferable to have a module that enforced some type of restriction on the individual performing the infraction. Examples could include but not be limited to preventing further file creates in home directory or preventing new jobs from being started.